Volume Graphics (2006)
T. Möller, R. Machiraju, T. Ertl, M. Chen (Editors)

Enabling Haptic Interaction with Volumetric MRI Data Through Knowledge-based Tissue Separation

Karljohan Lundin†, Claes Lundström‡, Matthew Cooper† and Anders Ynnerman†

† Norrköping Visualization and Interaction Studio, Linköping University, Sweden
‡ Center for Medical Image Science and Visualization, Linköping University, Sweden

Abstract

Direct volume haptics can provide both guidance and extra information during exploration of volumetric data. In this paper we present a novel approach to volume haptics enabling haptic exploration of tissue shape, borders and material properties in data despite low contrast and low signal to noise ratio, as is common in medical MRI data or low dose CT data. The method uses filtering based on implicit knowledge and addresses the problem of overlapping scalar ranges through the introduction of fuzzy classification and corresponding transfer functions for material properties as well as classification-based distance masking for haptic force direction.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Interaction techniques; H.5.2 [Information Interfaces and Presentation]: Haptic I/O; I.3.7 [Computer Graphics]: Virtual reality

1. Introduction

Direct volume rendering methods are becoming increasingly popular and new applications are found in a wide range of disciplines — one highly relevant example is pre-operative planning in which patient specific medical data, such as data from Computer Tomography (CT) and Magnetic Resonance Imaging (MRI) scans, is gathered. In the pre-operative planning work flow there is limited time for manual segmentation, pre-processing and calibration. The challenge is then to provide fast and accurate methods for direct visualization and interaction with patient specific data in a way that resembles the interaction with traditional medical models.

Research has shown that haptics has the potential to significantly increase both speed and accuracy of human-computer interaction [WH00, KD02, PNCD01, IN93, AKH01, WPS*02]. MRI data, however, is insufficient for producing convincing haptic feedback. In fact, to the authors' knowledge, there exists no method for direct haptic exploration of volumetric MRI data or other modalities with low contrast and signal to noise ratio. Poor contrast makes tissue shape estimation too uncertain and the overlapping tissue ranges of MRI renders direct material classification impossible. There is thus a need for more effective methods for producing the haptic feedback from these types of data.

In this paper we present an approach for providing haptic interaction with MRI data while avoiding time consuming manual segmentation. This is done by introducing a pre-processing step using knowledge-based tissue classification and surface information extraction. The developed methods are capable of dealing with noise as well as overlapping scalar tissue values and so produce high quality haptic rendering, guiding the user and presenting haptic cues about size and shape of target tissues, such as tumours.

2. Related Work

Direct volume haptics (DVH), follows two main approaches — force functions, introduced in [IN93], and constraints-based feedback, introduced in [LYG02]. Both these approaches have been combined into a single framework through the introduction of haptic primitives in [LGY05]. In [BG00] Bartz and Gürvit suggest a variation on the force function approach by generating the feedback from a dis-

† e-mail: {karlu,matco,andyn}@itn.liu.se
‡ e-mail: clalu@cmiv.liu.se

tance map instead of directly from the data. This distance map is used to produce a force that pushes the haptic probe away from the walls of a segmented cavity for navigation in MRI data.

The work presented in this paper builds upon the haptic primitives-based approach. The approach defines four haptic primitives, point, line, plane and force, representing basic haptic effects; shapes and forces. These primitives can be used to generate a wide range of haptic modes by simply controlling the primitive parameters strength, position and orientation as functions of the volumetric data. Surface feedback, used in this paper, is one of these modes.

We also employ the fuzzy classification method described in [LLY05]. The main characteristic of this method is that it translates straight forward domain knowledge of the data set into an efficient tissue separation measure. Examples of knowledge that can be exploited are tissue size and homogeneity, as well as proximity to other tissues. The classifying attributes used are *range weight* measures that describe the footprint of a given partial intensity range Φ in a voxel neighbourhood, in our work consisting of 13^3 voxels. The user provides range weight reference levels for each tissue corresponding to fully certain classification as this tissue. These reference levels are used to derive a *competitive classification certainty*, **P**, valid in the overlapping range and defined in the $[-1,1]$ interval, where the sign denotes the most likely tissue and the magnitude represents the classification confidence.

The MRI volume used in this project depicts a human liver with a cancer tumour and is of 10 bits precision and $256 \times 256 \times 128$ in resolution, see figure C1(i). We use two known facts that separate the tumour from other tissue in the same range: the tumour is relatively homogeneous and it is not surrounded by low-intensity regions such as air. Range weights corresponding to these two facts were retrieved, w_{r1} and w_{r2} for $\Phi_1 = [150,200]$ and $\Phi_2 = [0,120]$, respectively. The tumour references were set as $w_{r1} = 0.6, w_{r2} = 0.1$ and the other tissue references as $w_{r1} = 0.1, w_{r2} = 0.2$, resulting in the classification volume shown in figure C1(ii).

3. Enhanced Virtual Surfaces

Material information, such as strength and friction, and surface information, in our case the gradient vector, are usually both extracted from the scalar voxel data. Since MRI data often contains tissues with overlapping scalar ranges and noise that makes the use of the raw gradient impossible, we instead use separate methods for extracting material information and surface information.

3.1. Material Properties

To simplify the equations and descriptions, we will assume that we have three types of tissues (A, B and C) for which we want to control the haptic feedback and define different material properties. Two of the tissue types (A and B) have overlapping scalar ranges and can thus not be discriminated using a traditional transfer function, see figure 1(i).

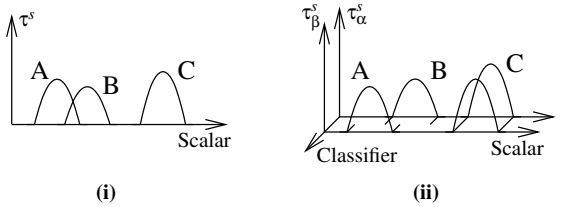

Figure 1: *Tissues with overlapping scalars (i) are separated in the classification dimension (ii).*

Two important material properties are extracted in this project, the strength of the haptic feedback and the friction of the rendered surfaces. These properties are specified through transfer functions, τ^s for strength and τ^μ for friction ($\tau : \mathcal{R} \rightarrow \mathcal{R}$). For tissues with overlapping scalar ranges, for example A and B in figure 1(i), the overlapping tissues are separated in a second dimension provided by the classification volume, C, derived from the value of **P** for each voxel in the MRI data [LLY05]. Thus, two separate transfer functions are defined for each material property, $\tau^s_\alpha/\tau^s_\beta$ and $\tau^\mu_\alpha/\tau^\mu_\beta$, see figure 1(ii), and the classification volume data is used to interpolate between the values of these two transfer functions. For example, the strength property, s, for a position, $\vec{x}$, is expressed as

$$s(\vec{x}) = \frac{1-C(\vec{x})}{2}\tau^s_\alpha(V(\vec{x})) + \frac{1+C(\vec{x})}{2}\tau^s_\beta(V(\vec{x})) \qquad (1)$$

where C is the classification volume and V is the MRI volume.

3.2. Surface Information

Since the unprocessed MRI data is insufficient for use as surface information, this information must, instead, be derived from the tissue distribution information. We do this by generating a separate volume containing a soft gradient near the tissue borders. First tissue masks, M_i, are generated for each tissue ($i \in \{\mathrm{A},\mathrm{B},\mathrm{C}\}$), which can be viewed as an implicit segmentation. For this a separate strength transfer function pair for each single tissue type is first defined: $\tau^s_{\alpha,i}$ and $\tau^s_{\beta,i}$. To produce the binary mask, the transfer function pair is applied to the dataset and thresholded,

$$M_i = \begin{cases} 1, & \frac{1-C(\vec{x})}{2}\tau^s_{\alpha,i}(V(\vec{x})) + \frac{1+C(\vec{x})}{2}\tau^s_{\beta,i}(V(\vec{x})) \geq T \\ 0, & \frac{1-C(\vec{x})}{2}\tau^s_{\alpha,i}(V(\vec{x})) + \frac{1+C(\vec{x})}{2}\tau^s_{\beta,i}(V(\vec{x})) < T \end{cases} \qquad (2)$$

where T is the threshold chosen for the masking, generally chosen to be lower than half of the typical tissue strength specified in τ^s.

The mask volume, M_i, can now be used to generate a volume with explicit surface information for the tissue type. This is done by applying the distance operator to the mask volume, to produce a volume where every voxel earlier classified as part of the tissue type will contain a value describing its distance to the closest non-tissue voxel. This will provide a soft gradient inside the tissue volume, suitable for haptic rendering, see figure C4 (top). The data obtained is similar to that used by Bartz et al [BG00] to push the haptic instrument away from the walls of a cavity, however the extraction procedure, use and subsequent haptic effect differs in our work.

Since the distance maps of the separate tissue types are spatially non-overlapping, their surface information can be merged into a single volume through voxel-wise addition. This principle is shown in figure C4.

3.3. Haptic Surfaces

Our implementation uses two haptic primitives: a plane primitive provides surface feedback and a line primitive simulates friction. Both the plane and line primitives have an orientation, $\vec{q}$, being the normalized gradient vector of the surface information volume, D, as long as the gradient is non-zero. If not, an arbitrary vector is used and the primitives' strengths are set to zero so that no feedback is provided. For a more detailed description of how the haptic rendering of virtual surfaces is implemented, we refer to [LGY05].

The plane primitive strength is defined by equation 1. The line primitive strength, s_μ, that is the friction feedback, is evaluated as the friction value, μ, multiplied by the estimated surface force. The surface force is estimated to be whichever is the smaller of the surface strength, s, and the normal directed feedback force,

$$s_\mu = \mu \min(s(\vec{x}_{\text{proxy}}), k\vec{q} \cdot (\vec{x}_{\text{proxy}} - \vec{x}_{\text{probe}})) \qquad (3)$$

where k is the stiffness of the decoupling used when solving the haptic primitives, μ is estimated identically to s in equation 1, and $\vec{x}_{\text{proxy}}$ and $\vec{x}_{\text{probe}}$ are the proxy and probe positions.

3.4. Noise Reduction

Sponge-like tissues, such as the liver, have natural small holes that do not provide constructive contribution to the haptic feedback, but which disturb the distance function. Assuming that these holes are generally small in size and unconnected to each other, removal can be performed using standard 3D morphological transformations as described in, for example, [SHB99]. We perform the closing operation, to remove minor holes and cracks, and then an opening operation, to remove unevenness in the mask surfaces. The results from the morphological operations are shown in figure C2.

4. Haptic Margins

A common task in pre-operative planning is finding measures for the operation target. In these cases a user may be more effectively assisted by haptic support at a distance from the target tissue than, as is described above, at the tissue borders. This has the potential to guide the user to find an appropriate margin while measuring, for example, a cancer tumour.

Our method puts us in a position to implement haptic margins through a straightforward addition of an alternative distance map containing, for each voxel *not* part of the target tissues, the distance to the closest *target tissue* voxel. This distance map represents the surface information of the tissues at a distance from the real tissue surface. Since the haptic interaction will be outside the actual tissue, however, the scalar value of the MRI volume can no longer be used to control the material properties as described in section 3.1. The simplest approach is to use constant material properties, for example properties representative of the current type of tissue or any properties suitable for the physical support needed for accurate measurements. The strength property, s, can also be modulated using the value from the distance map, providing distance specific strength.

5. Implementation

Our application is implemented using the H3D API and the Volume Haptics Toolkit [LCP*06] (VHTK). H3D is an X3D-based system for implementing multi-modal applications and VHTK extends the functionality of H3D by providing scene-graph nodes for generating haptic feedback from volumetric data, and volume data handling, filtering and visualization. The approach presented in this paper is implemented as multiple, configurable filters for the toolkit. Also, the haptic mode for rendering surface feedback from scalar data has been modified to allow for the tissue separation. Selection between the two methods described in sections 3 and 4 is made by replacing the surface information volume and the transfer functions provided for the haptic mode described in section 3.3. For the visual rendering we deploy a GPU accelerated, proxy geometry volume renderer. This modified renderer accepts a pair of RGBA transfer functions that are, in a pre processing step, converted into a look-up table, see figure C3, that is used by a simple fragment shader to obtain the fragment colour.

The filters, visual rendering and haptic schemes are fitted into the multi modal scene-graph provided by H3D. The scene-graph used to render the methods presented in section 3 are shown in figure 2.

6. Results

Using our approach for haptic surface information and tissue separation, we obtain a distinct and clear haptic feedback from the different tissues. Providing haptic feedback

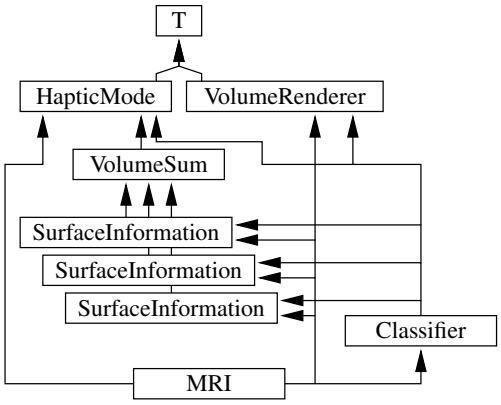

Figure 2: *The scene-graph for the experimental setup. T represents the transform that control the location of the haptic and graphic representations of the MRI data.*

from the tumour tissues without using our tissue separation would produce severe problems with haptic artifacts and feedback from surrounding tissues, as is visually indicated in figure C5. Even though the MRI data is noisy and has contrast generally unusable for haptic feedback, our approach makes it easy to find and follow the tissue surfaces through haptic probing. The surface properties can also be fully controlled through the provided haptic transfer functions. If the applied force exceeds the strength specified in the transfer function, the haptic probe penetrates the implicit surface. In this way the user can push through tissues surrounding the cancer tumour, thereby avoiding occlusion.

The haptic margin rendering provides a distinct support at any distance from the surface of the target tissue and while tracing the tissue contour, the feedback provides a sense of the tumour shape. By applying a pressure to the haptic probe, the outer parts of the margin shell can be penetrated so that measures can be taken at a smaller distance from the tissue surface.

7. Conclusions

We have presented a novel method enabling haptic interaction with MRI data, thereby providing physical guidance and extra information about material and shape through the sense of touch. The approach is based on domain knowledge filtering of the MRI data to extract separate material and surface information and so avoids time consuming explicit classification. This solves the problem with overlapping scalar ranges and circumvents the problem with low contrast and low signal to noise ratio. Even though this approach has so far only been tested on MRI data, we anticipate that it will be useful also for interaction with other problematic modalities, such as low dose CT.

References

[AKH01] ADAMS R. J., KLOWDEN D., HANNAFORD B.: Virtual training for manual assembly task. *Haptics-e, The Electronic Journal of Haptics Research (www.haptics-e.org) 2*, 2 (October 2001).

[BG00] BARTZ D., GÜRVIT Ö.: Haptic navigation in volumetric datasets. In *Proceedings of PHANToM User Research Symposium* (2000).

[IN93] IWATA H., NOMA H.: Volume haptization. In *Proceedings of IEEE 1993 Symposium on Research Frontiers in Virtual Reality* (October 1993), pp. 16–23.

[KD02] KIRKPATRICK A. E., DOUGLAS S. A.: Application-based evaluation of haptic interfaces. In *Proceedings of the 10th Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems* (2002).

[LCP*06] LUNDIN K., COOPER M., PERSSON A., EVESTEDT D., YNNERMAN A.: Enabling design and interactive selection of haptic modes. *Virtual Reality Journal* (2006). (to appear).

[LGY05] LUNDIN K., GUDMUNDSSON B., YNNERMAN A.: General proxy-based haptics for volume visualization. In *Proceedings of the World Haptics Conference* (Pisa, Italy, March 2005), IEEE, pp. 557–560.

[LLY05] LUNDSTRÖM C., LJUNG P., YNNERMAN A.: Extending and simplifying transfer function design in medical volume rendering using local histograms. In *Proceedings of Eurographics/IEEE VGTC Symposium on Visualization* (2005).

[LYG02] LUNDIN K., YNNERMAN A., GUDMUNDSSON B.: Proxy-based haptic feedback from volumetric density data. In *Proceedings of Eurohaptics* (2002), University of Edinburgh, United Kingdom, pp. 104–109.

[PNCD01] PASSMORE P. J., NIELSEN C. F., COSH W. J., DARZI A.: Effects of viewing and orientation on path following in a medical teleoperation environment. In *Proceedings of IEEE Virtual Reality 2001* (2001).

[SHB99] SONKA M., HLAVAC V., BOYLE R.: *Image Processing, Analysis, and Machine Vision*, 2 ed. PWS Publishing, 1999.

[WH00] WALL S., HARWIN W.: Quantification of the effects of haptic feedback during a motor skills task in a simulated environment. In *Proceedings at Phantom User Research Symposium'00* (2000).

[WPS*02] WALL S. A., PAYNTER K., SHILLITO A. M., WRIGHT M., SCALI S.: The effect of haptic feedback and stereo graphics in a 3d target acquisition task. In *Proceedings of Eurohaptics* (2002), University of Edinburgh, United Kingdom.